Texture Modelling with Nested High–order Markov–Gibbs Random Fields ☆

Ralph Versteegen, Georgy Gimel'farb, Patricia Riddle

Department of Computer Science, The University of Auckland, Auckland 1142, New Zealand

arXiv:1510.02364v1 [cs.CV] 8 Oct 2015

Abstract

Currently, Markov–Gibbs random field (MGRF) image models which include high-order interactions are almost always built by modelling responses of a stack of local linear filters. Actual interaction structure is specified implicitly by the filter coefficients. In contrast, we learn an explicit high-order MGRF structure by considering the learning process in terms of general exponential family distributions nested over base models, so that potentials added later can build on previous ones. We relatively rapidly add new features by skipping over the costly optimisation of parameters.

We introduce the use of local binary patterns as features in MGRF texture models, and generalise them by learning offsets to the surrounding pixels. These prove effective as high-order features, and are fast to compute. Several schemes for selecting high-order features by composition or search of a small subclass are compared. Additionally we present a simple modification of the maximum likelihood as a texture modelling-specific objective function which aims to improve generalisation by local windowing of statistics.

The proposed method was experimentally evaluated by learning high-order MGRF models for a broad selection of complex textures and then performing texture synthesis, and succeeded on much of the continuum from stochastic through irregularly structured to near-regular textures. Learning interaction structure is very beneficial for textures with large-scale structure, although those with complex irregular structure still provide difficulties. The texture models were also quantitatively evaluated on two tasks and found to be competitive with other works: grading of synthesised textures by a panel of observers; and comparison against several recent MGRF models by evaluation on a constrained inpainting task.

Keywords: texture synthesis and analysis; high-order MRFs; local binary patterns; structure learning.

1. Introduction

Texture modelling is central or important to many computer vision and image processing tasks such as image segmentation, inpainting, texture classification or synthesis, anomaly (defect) detection, and image recognition. Although successful specialised algorithms for texture classification, synthesis and segmentation have been developed, generative probabilistic models which offer relatively complete models of statistics of individual textures are appealling. They may be applied not only to all of the above tasks, but to any where appearance priors or feature extraction are needed, and they are also of interest to understanding human vision. Generative models must capture most of the features of a texture that are significant to human perception in order to be successful, whereas texture features used for discrimination need not.

The most prevalent tool for image and texture modelling are Markovian undirected graphical models, a.k.a. Markov random fields (MRFs). An MRF together with an explicit Gibbs probability distribution (GPD) is called herein a Markov–Gibbs random field (MGRF). MGRFs are particularly popular for image analysis involving the determination of boundaries (as in segmentation) or enforcing smoothness (e.g. in stereoscopic matching and image denoising). In these cases the Markov networks are usually sparse, with the directly interacting neighbours of each variable being close by. Some high-order MGRF models have been proposed for such tasks (e.g. [1–3]), and for binary variables efficient maximum a posteriori (MAP) algorithms exist, such as graph cuts [4]. However things are different in the domain of image and texture modelling, where inference needs to be performed on real-valued or highly multivalued image variables in dense Markov networks. The networks used in this paper typically have Markov blankets containing 50–100 nearby and distant pixels, and even sampling from the models proves to be difficult.

MGRFs and other probabilistic texture models reduce images $\mathbf{g}$ to a vector of statistics of image features $\mathbf{f}(\mathbf{g})$, which are assumed sufficient to describe the texture. The model is completed by assigning an energy ϕ to each feature vector, giving a Gibbs probability distribution over images:

$$p(\mathbf{g}) \propto \exp(-\phi(\mathbf{f}(\mathbf{g}))). \tag{1}$$

Historically statistics of pairs of pixels [5–7] were used. However higher-order MGRFs (which cannot be expressed in terms of lower order ones), have become more common as they are

☆©2015. Licensed under Creative Commons CC-BY-NC-ND 4.0
http://creativecommons.org/licenses/by-nc-nd/4.0/

Email addresses: rver017@aucklanduni.ac.nz (Ralph Versteegen), g.gimelfarb@auckland.ac.nz (Georgy Gimel'farb), p.riddle@auckland.ac.nz (Patricia Riddle)

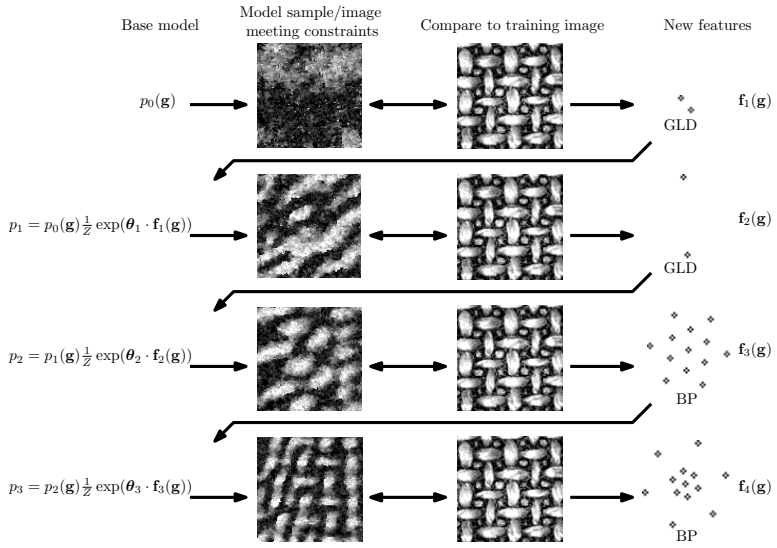

Figure 1: Nesting exponential-family texture models. At each nesting iteration images meeting the existing statistical constraints of the current model are generated, such as by sampling from the model. One or more new features/potentials (in this example grey level differences (GLD) and binary patterns (BP)) are selected at each iteration by searching for the largest deviations of their empirical marginal distributions against the target image. Each provides a new set of constraints, and adds a Gibbs factor to the model, thereby moving the model closer to the target.

recognised to be necessary for more expressive models of natural images and textures (e.g. [2, 3, 8–11]). Higher order interactions in image models allow for abstracting beyond pixels, building upon larger scale image attributes like edges, and for context and complex structures to be captured. In addition, since regularly tiled textures have strong long range correlations between nearby tiles it is natural to learn an interaction structure (i.e. the pattern of statistic dependences between pixels) specific to the texture. Yet it is still almost unheard of in computer vision and image modelling for higher-order MRF structure to be learned rather than hand selected.

However, selection of high-order features poses significant problems. The cardinality of a space of possible feature functions grows combinatorially in the order, due to both freedom in the shape of the support (variables/pixels to select as input), and the need to reduce or manage its high dimensional input domain. In other words the features should be parameterised with a reasonable number of parameters. The higher-order MGRFs in use nearly exclusively apply linear filters as feature functions, with statistics of the filter responses, such as means and variances [12], correlations [13], or histograms [14] forming a description vector. For texture classification many other methods of extracting useful information out of a high dimensional pixel co-occurrence matrix have been investigated (e.g. [15, 16]). Dimensionality can be reduced by making assumptions such as that images are invariant to contrast and offset changes. However this approach has seemingly been little-applied to generative texture models.

In order to tackle these problems, we build texture models by a model nesting procedure which greedily selects features and can build higher order features by composing lower order ones. Unlike some other works (e.g. [13, 14]) we do not attempt to provide a fixed set of statistics/texture features to distinguish between all textures (a goal with the Julesz conjecture [17] as its origin), but rather learn texture-specific features. This potentially provides compact representations while still allowing a large and varied space of descriptors. Each nesting iteration corrects statistical differences between the training image and the textures class given by the previous model, as sketched in Figure 1.

Contributions of this paper are as follows: (*i*) We efficiently select high-order features by "nesting" models with heterogeneous features/potentials, while coping with the difficulties of inference in dense MGRF texture models. Unlike the model nesting used previously in [14, 18] we do not learn MLE parameters at each nesting iteration, which is very expensive, but instead generate images which match the current statistical constraints (Section 4.3). These are equivalent to samples from the ideal MLE model. Correct parameter learning can be delayed until afterwards. We use no hidden variables as is currently popular which eases learning and inference, with parameter learning remaining convex in theory. (*ii*) We extend the very popular local binary pattern (LBP) descriptors of images by learning the offsets of the surrounding pixels (Section 4.6) for use as high-order 'binary pattern' (BP) MGRF texture features. These are quite different from the common high-order linear filtering or Potts potentials, and faster to compute than responses of large linear filters. LBPs have apparently never been used in this way despite enormous popularity as image descriptors. Experiments into texture synthesis using MGRFs with LBP histograms as sufficient statistics can provide insight into the visual features actually captured by the LBPs. (*iii*) We compare several families of nested texture models utilising different high-order features, including different methods of selecting BP offsets. The resulting texture models have heterogeneous feature sets composed of second-order grey level difference (GLD) features, and of up to 13^{th}-order BP features or Laplacian of Gaussian and Gabor filters. The use of learned long range GLD interactions

allows almost-regular (tiled) textures in particular to usually be synthesised well. (*iv*) The ability of the proposed procedure to learn characteristic features across different types of textures is demonstrated with texture synthesis across a varied set of greyscale textures, also evaluated by a panel of observers, along with several other comparisons. (*v*) In order to improve generalisation and to attempt to allow partially inhomogeneous training images a variant on the maximum likelihood estimate (MLE) was used, such that the training image is split into pieces and the minimum of the likelihoods of the pieces rather than their product is maximised.

2. Related Work

2.1. High-order MGRF models

Much research in texture analysis has focused on describing textures using the distributions of responses of square linear filters. MGRF texture models utilising non-trivial filters were introduced with FRAME [10, 14], where the filters were selected from a manually-specified bank. More recently learning the filters themselves (with predetermined fixed supports) has been popularised by the Field-of-Experts (FoE) model [11]. These model the marginal distibution of responses for each filter with hand-picked potentials with few parameters; in the original FoE model they are student-t distributions. Therefore the interaction structure is learned as filter coefficients. FoE was extended to bimodal FoE (BiFoE) which uses more informative bimodal potentials, and successfully applied to texture modelling by Heess et al. [9]; several state of the art generative texture models have been built on BiFoE, some using various configurations of hidden variables. Kivinen and Williams [19] improved on BiFoE by using gated MRFs [20], and Luo et al. [21] investigated convolutional deep belief networks (DBN) and spike-and-slab potential functions. However because of learning difficulties and to reduce required computation all these learned-filter models have been restricted to relatively small filter sizes. These do not directly capture distant interactions; the largest filter size used in the mentioned works was 11×11. As a result, while these filter-based MGRFs model certain classes of textures excellently, they have inherent weaknesses. A more general survey of MRFs which covers high order models can be found in [3].

The recently popular hierarchical filter-based MRFs are a merger of MRF models, especially Boltzmann machines, and models of single image patches using independent components analysis [22] and overcomplete Product-of-Experts [23]. These are applied to larger images by tiling and overlapping the filters. Today MRFs including latent variables are popular, often inspired by simple and complex cells in the visual cortex, with strong links to convolutional artificial neural networks. Using multiple layers allows application to high level computer vision tasks where the Markov property does not hold, such as object recognition and face modelling. In many cases integrating out the latent variables leads to a completely connected, non-Markovian interaction graph. We diverge from this direction to consider the direct inclusion of higher-order features in models, which is more common for other machine learning tasks where there often is no analogue to linear filtering.

2.2. Texture features

Various high-order local texture descriptors have been applied to texture classification (e.g. [15, 16, 24]) going back at least 20 years. For such applications, invariances to contrast, offset, scale, rotation, and deformation are given particular weight. Non-linear texture features have also been used in MGRF texture models; Sivakumar and Goutsias [25] introduced MGRF texture models which used sophisticated multiscale features defined through mathematical morphology. Many sought to directly reduce the dimensionality of high-order co-occurrence histograms through traditional techniques such as spectral clustering of histogram bins [26], vector quantisation [16], Gaussian mixture models [24], and self-organising maps [15]. However such dimensionality reduction techniques usually require a nearest-neighbour search to map new data points, which likely makes them unsuitably slow as feature functions in MGRFs.

Probably the most popular texture descriptors for discrimination are LBPs [27], which compare the intensities of a ring of pixels $p_1, \ldots, p_k$ around a central one p_0, and form the bit-vector $(p_i > p_0)_{1 \le i \le k}$. These are nearly contrast- and offset-invariant and can be straight-forwardly extended to rotation invariance [28] by merging bins, and extended to partial scale invariance by using multiple concentric rings. In addition to their proven ability to distinguish textures LBPs are also very cheap to compute, hence we investigate their use in texture modelling. Nosaka et al. [29] introduced co-occurrence statistics of neighbouring LBPs in a way that is rotationally invariant, This is particularly interesting for future extension of the BP-based MGRFs in this paper to rotational invariance. LBPs have inspired a number of other variants such as local ternary patterns (LTPs) [30], and local radial index [31]. Liu et al. [32, 33] have also used LBPs, LTPs, and related non-linear "ordinal" texture descriptors with learnt offsets, and applied them to texture classification and retrieval. Although explained within the framework of MGRFs, these approaches do not involve learning a complete MGRF model. They built up the higher-order cliques (up to 20^{th} order) out of the lower-order ones by searching for cliques in an interaction graph computed using a thresholding rule. We find this too restrictive and pursue alternative strategies for selecting the offsets.

2.3. Structure learning and nesting

A method for selecting the features of the model which is both intuitive and theoretically sound is to use greedy *sequential structure selection*, as has been used by various authors [14, 18, 34–36] with a number of details varying. This alternates between adding one or more features/factors to the current model from a candidate set according to an estimate of the best feature to add, and then finding the new MLE of the parameters (initialised at those from the previous iteration). The most common metric for selecting the best feature to introduce is to select that with the largest 'error' between training image

and model synthesis result [14, 34, 35]. This paper follows the same general approach, which we call 'nesting', although we describe our particular flavour. The main difference in our approach is to attempt to directly generate images matching the statistics at each iteration, learning approximate parameters as a side effect.

Della Pietra et al. [18] treated structure learning in MGRFs (for natural language processing problems) as a feature selection problem, and built up higher order features out of lower order ones. Several other authors have since considered selecting MGRFs features by gradually composing together low order atomic features ('general-to-specific') (e.g. [37, 38]), or by starting from template-like features (sometimes called 'patterns') which are conjunctions of simple $x_i = c_j$ predicates, derived directly from the training data and gradually generalising them ('specific-to-general') (e.g. [39]).

A closely related popular method for MGRF structure learning is the use of sparsity-inducing L_1 regularisation [35, 40], which forces some feature weights to exactly zero so that they can be removed. Otherwise proceeding like nesting, this approach has the advantages that the regularisation combats overfitting, that it allows removing a feature after it has been added, and is a convex problem. Recently Chen and Welling [41] suggested the use of spike-and-slab priors (similar to L_0 regularisation) instead of L_1 regularisation. This has some advantages, but does not also lead to an efficient greedy algorithm with an optimality guarantee.

2.4. Texture synthesis

Practical texture synthesis is currently dominated by algorithms which combine pieces of a source image so that the pieces fit together well. This is defined in terms of the match between neighbourhoods of the pieces. Region-growing techniques add one pixel (e.g. [42, 43]) or patch (e.g. [44, 45]) to the image at a time. Algorithms are often multiscale [43, 46, 47]. However all these synthesis algorithms have the property that they copy large parts of the source image directly into the result, either an explicit part of the algorithm, or an emergent property of the algorithms' search for best matching neighbourhoods, due to 'forced moves'.

Another rather successful but not realtime class of texture synthesis algorithm is based on global optimisation of the image such as by projection onto the set of images with certain statistics equal to those of the training image, in particular statistics of wavelet-decompositions [12, 13]. A codebook of texture patches or primitives can also be used instead [48]. These are implicitly MRF texture models as they provide no explicit probability distributions (hence our discrimination between MGRFs and MRFs) and so are less broadly applicable than probabilistic models. Hence, although these give good synthesis results these algorithms in fact attack a different problem than the more statistical one that this paper does; we use texture synthesis only for evaluation.

3. Nested Markov–Gibbs Random Fields

This section provides first some definitions, and then defines general exponential distributions as the conceptual and theoretical basis for MGRF model nesting. A description of the nesting algorithm follows.

3.1. Basic notation and definitions

Any probability distribution which is nowhere-zero can be represented as a GPD factorised into Gibbs factors: functions of complete subgraphs (cliques) of the Markov network. In this paper we consider the class of MGRFs with factors (synonymously, potentials) that can be described as the product of a fixed feature function —identifying a certain signal configuration (pattern)— and a corresponding weight/parameter.

We restrict our scope to modelling of homogeneous textures by repeating the Gibbs factors and cliques across the image to achieve translation invariance. Let $\mathcal{R} \subset \mathbb{Z}^2$, where $\mathbb{Z}$ are the integers, be a finite lattice of image coordinates and $\boldsymbol{\alpha} = \left\{\alpha_i : \alpha_i \in \mathbb{Z}^2; 1 \leq i \leq d\right\}$ be a list of d coordinate offsets with $\alpha_1 = (0, 0)$ fixed. An order d clique family $C_{\boldsymbol{\alpha}}$ in $\mathcal{R}$ is the set of all spatially repeated cliques with the offset pattern given by $\boldsymbol{\alpha}$:

$$C_{\boldsymbol{\alpha}} := \{(r_1, \ldots, r_d) : r_1, \ldots, r_d \in \mathcal{R};\ r_i - r_1 = \alpha_i;\ i = 1, \ldots, d\}.$$

Let $\mathbf{g} : \mathcal{R} \to \{0, \ldots, Q-1\}$ be an image on $\mathcal{R}$ with Q possible grey levels. An order d feature is a function $f : \{0, \ldots, Q-1\}^d \to \mathbb{N}_s$ with finite range $\mathbb{N}_s := \{1, \ldots, s\}$ and an associated clique family C_f given by offset list $\boldsymbol{\alpha}_f$. The histogram of values of f collected over an image $\mathbf{g}$ is the vector

$$\mathbf{h}_f(\mathbf{g}) = \left[\sum_{c \in C_f} [\![l = f(\mathbf{g}_c)]\!] \ :\ l \in \mathbb{N}_s\right] \tag{2}$$

where $[\![\cdot]\!]$ is the Iverson bracket mapping true $\mapsto 1$, false $\mapsto 0$, and $\mathbf{g}_c$ is notation for sequence of pixels $\mathbf{g}(r_1), \ldots, \mathbf{g}(r_d)$, $r_i \in c$, c a clique. The order of a model is defined as the maximum order of any of its features.

3.2. General exponential distributions

In the nesting procedure, new features to be added to the current model are selected based on the disagreement between their statistics in the training image and the model's expected statistics, approximated from model samples. In this way the model is modified to correct these errors. When the statistics in question are expectations (equivalently, histograms), these corrections take a particularly simple form.

Let F_{i+1} be a set of feature functions and let the histogram vector $\mathbf{f}_{i+1}(\mathbf{g})$ be the concatenation of histograms for the set F_{i+1}: $\mathbf{f}_{i+1}(\mathbf{g}) := [\mathbf{h}_f(\mathbf{g}) : f \in F_{i+1}]$. A general exponential family model [49] is a probability distribution which can be written in the following form specified by a base model $p_i(\mathbf{g})$, a vector-valued feature function $\mathbf{f}_{i+1}(\mathbf{g})$, and a parameter vector $\boldsymbol{\theta}_{i+1}$,

$$p_{i+1}(\mathbf{g}|\boldsymbol{\theta}_{i+1}) = \frac{1}{Z(\boldsymbol{\theta}_{i+1})} p_i(\mathbf{g}) \exp(-\mathbf{f}_{i+1}(\mathbf{g}) \cdot \boldsymbol{\theta}_{i+1}) \tag{3}$$

where $Z(\boldsymbol{\theta})$ is a normalisation factor, and $\cdot$ denotes dot-product. In our case, the parameters are a concatenation of per-feature parameter vectors $\boldsymbol{\theta}_{i+1} = [\boldsymbol{\theta}_f : f \in F_{i+1}]$.

If one wishes to find a model p_{i+1} meeting constraints to be satisfied in the form of expectations $\mathbb{E}_{p_{i+1}}[\mathbf{f}_{i+1}(\mathbf{g})] = \mathbf{f}_{i+1}(\mathbf{g}^\circ)$ (where $\mathbf{g}^\circ$ is given training data), but already has prior information expressed as a base model (in this case the model at the previous iteration, p_i), then it is widely known [18] that the probability distribution p_{i+1} which matches the new constraints but deviates from the base model the minimum possible amount (as measured by the Kullback-Leibler divergence, i.e. has maximum entropy relative to p_i) has a simple form. It is the general exponential family model given in (3) with the maximum likelihood estimate (MLE) of the parameters $\boldsymbol{\theta}^*_{i+1} := \arg\max_{\boldsymbol{\theta}} p_{i+1}(\mathbf{g}^\circ|\boldsymbol{\theta})$. It can be seen from Eq. (9) that the MLE parameters achieve the desired statistical constraints. Technically, the MLE $\boldsymbol{\theta}^*_{i+1}$ may lie at an infinitely distant point in parameter space, which is not permitted by the usual definition of a GPD. This possibility is easily avoided by slightly smoothing $\mathbf{f}_{i+1}(\mathbf{g}^\circ)$.

3.3. Model nesting

The challenges of both selecting high-order features and structure can be met by gradually building up features from pieces, reducing the set of candidate features to a size that can be exhaustively searched at each iteration. One could reasonably expect that when some configuration of pixels (as recognised by a feature function) is characteristically common for a texture, that the configuration restricted to a subset of pixels would usually also be common. Hence we conjecture that in practice high-order feature functions picking out characteristic interactions can be found by building up from lower order features.

From the theory of general exponential distributions we can see that there is no need to modify whatever parameters the base model may have, although in most previous works involving sequential structure learning this is normally done as it allows increasing the likelihood. However to attempt to reduce overfitting we suggest that it is reasonable to only adjust parameters associated with the new features $\mathbf{f}_{i+1}$. Dudík et al. [50] described such a coordinate-wise descent algorithm for a wide class of generalised maximum-entropy problems which selects a best single parameter to update on each iteration. They state the standard approach of updating all parameters each iteration "is impractical when the number of features is very large." However in our experiments we found that this was too aggressive a solution to overfitting, and obtained better results by keeping all parameters free.

Given a set of feature functions F_i, having already been selected, a *selector* function $C(F_i)$ is defined to provide a candidate set of new features, possibly built upon previous features. The algorithm proceeds stage-wise through a sequence of selectors $C^1, \ldots, C^k$ each providing features of a certain order and type, to avoid the problem of comparing across feature types. Sequential selection simply selects one or more feature functions among $C^j(F_i)$ on which the current model p_i has the largest disagreement (error) against the training data $e(f) := d(\mathbf{h}_f(\mathbf{g}^\circ), \mathbf{h}_f(\mathbf{g}^{(i)}))$, where d is a distance function on histograms. The next selector in the sequence is proceeded to after $\max_{f \in C^j(F_i)} d(\mathbf{h}_f(\mathbf{g}^\circ), \mathbf{h}_f(\mathbf{g}^{(i)}))$ becomes too small, and the algorithm terminates when the selectors are exhausted.

One possible variant is to use separate training and validation datasets, and to stop when performance (measured either by d or an application-specific external evaluation function such as texture discrimination accuracy) on the validation begins to decrease. Such a validation test was used in the earlier work [51], in order to detect slightly non-homogeneous textures. This paper takes the idea further that statistical variations across the training data are important: Section 4.1 discusses the selection of features from a set of pieces of the training images rather than less powerfully only looking for variation between a single training and validation pair.

Previous sequential feature selection implementations have used ℓ_1 [14, 40] or ℓ_2 distance or 'gain' for d. Gain is the increase in information contained in a model (equivalently, decrease in entropy) due to adding a feature to it, and can be either be estimated analytically [10] or, at great expense, by actually comparing every possible extended model [18]. Zhu et al. [10] also gave an correction to the gain for theoretical uncertainty in the estimated statistics, which is a form of the Akaike information criterion (AIC), penalising model complexity.

Algorithm 1 Outline of basic MGRF nesting procedure

Require: Initial base MGRF model p_0 with feature set F_0, training image $\mathbf{g}^\circ$
  $i \leftarrow 0$
  for each selector C in $C^1, \ldots, C^k$ **do**
    loop
      Obtain set of samples $\{\mathbf{g}_l\}_{1 \le l \le n}$ from model p_i
      for each $f \in C(F_i)$ **do**
        Collect histograms $\mathbf{h}_f(\mathbf{g}^\circ)$ and $\mathbf{h}_f(\mathbf{g}_1), \ldots, \mathbf{h}_f(\mathbf{g}_n)$
        $\bar{\mathbf{h}}_f(\bar{\mathbf{g}}) \leftarrow \frac{1}{n} \sum_l \mathbf{h}_f(\mathbf{g}_l)$
      end for
      Pick one or more f with maximal $d(\bar{\mathbf{h}}_f(\bar{\mathbf{g}}), \mathbf{h}_f(\mathbf{g}^\circ))$ for some histogram distance measure d
      $F_{i+1} \leftarrow F_i \cup \{f\}$
      Form p_{i+1} by adding f to p_i as in Eq. (3)
      Learn parameters $\boldsymbol{\theta}_{i+1}$ for the new features
      $i \leftarrow i + 1$
    end loop
  end for

However, real uncertainty in estimates found through Markov chain Monte Carlo (MCMC) sampling from a MGRF is typically vastly larger than theoretical estimates of uncertainty because the sampling may not converge (see Section 4.3). This is especially true because the quality, and also statistics, of the approximate samples that we attempt to rapidly draw from the current texture model are highly variable.

We used the Jensen–Shannon divergence (JSD) [52] for distance function d as a proxy to the additional information content

of $\mathbf{h}_f(\mathbf{g}^\circ)$. The JSD is defined as

$$D_{\mathrm{JS}}(p \parallel q) := \frac{1}{2}\left(D_{\mathrm{KL}}\left(p \parallel m\right) + D_{\mathrm{KL}}\left(q \parallel m\right)\right)$$

where D_{KL} is the Kullback-Leibler divergence (KLD), p and q are probability distributions, and m is the averaged distribution $\frac{1}{2}(p+q)$. The JSD measures the discriminability of two distributions, namely the average certainty (as measured in bits) with which a sample can be ascribed to one or the other. The JSD is symmetric, bounded in the range $[0, 1]$ bits (when computed using base 2 logarithms) and is popular because it is suitable for comparing two empirically estimated distributions, unlike the KLD, which is often undefined. In practice, the JSD usually provides very similar rankings to the ℓ_1 distance.

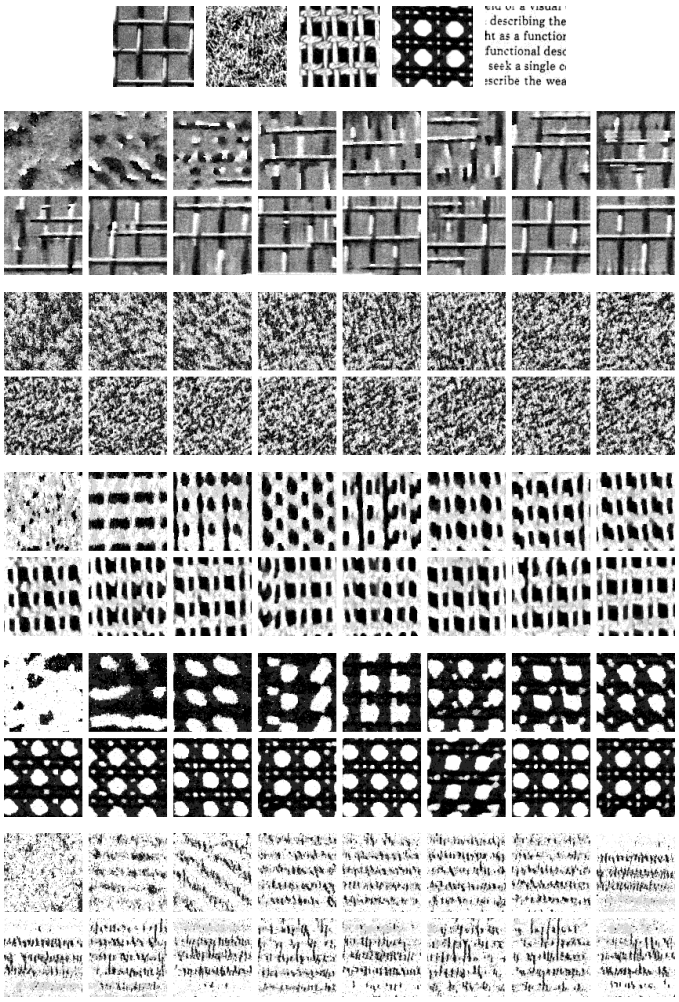

Figure 2: Examples of the nesting procedure. The top row shows pieces of the original five textures (from left to right: D1, D9, D20 and D101 [53] and 'text' [13]). The subsequent five pairs of the rows illustrate the progress of feature selection. Each image (from left to right) shows one of the CSA 'samples' at a nesting iteration, which approximately matches the statistical constraints of the current model (see Section 4.3). The first row of each pair shows addition of pairwise GLD features, and the second shows addition of the 9$^{\mathrm{th}}$-order jagged star features (see Section 4.6).

Algorithm 1 summarises this approach. Figure 2 gives examples of its operation, starting from the initial model (from which the first sample is drawn) which included only marginal and nearest-neighbour pairwise potentials and proceeding through two selectors. The first selector returns three pairwise potentials at a time. Frequently, there is no visible improvement from one step to another, however even the samples getting worse do not indicate that the model is getting worse, due to the unstable sampling process. There is also usually a big jump as each new class of features is added, and decreasing returns thereafter. Determining a suitable stopping point for model learning is difficult, as the histogram distances $e(f)$ do not reflect visual proximity to the training image very well, and a texture similarity metric such as improved structural texture similarity (STSIM2) [55] might be a better option. For experiments in this paper we instead used a fixed 8 iterations for each selector, as this removes the variability due to an imperfect stopping rule, which can have a large effect. Figure 3 shows the cliques selected in two of the examples in detail.

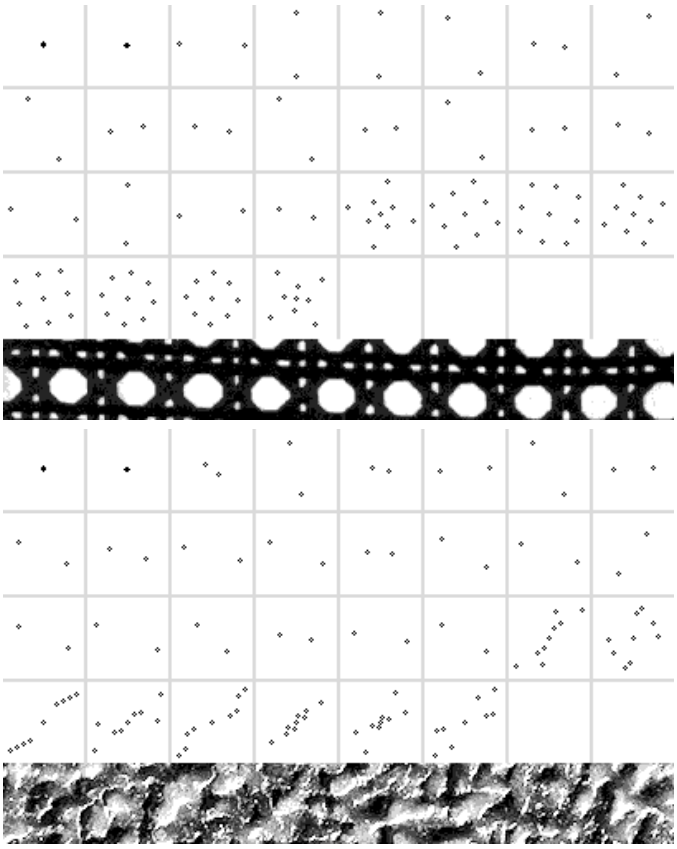

Figure 3: Shapes of the 2$^{\mathrm{nd}}$- and 9$^{\mathrm{th}}$-order 'jaggered star' cliques selected for D101 and Metal0004 [54], in order of selection. A piece of the texture at the same scale is shown at bottom.

4. Texture Modelling

The procedure described above is generic and could equally well be applied to learning MGRF models for machine learning tasks outside of computer vision. Its refinement specific to texture-modelling, including some of our implementation details is given below.

Broadly, there are two ways to achieve desired invariances: by designing the models to enforce those invariances (e.g. by using feature functions with some invariance property), or by learning them from suitable training data which demonstrates them. We use both approaches, although the latter is only preliminarily developed here.

4.1. Texture-specific maximisation of the minimum likelihood

Up to now, we have described model learning in terms of MLE that maximises the total probability of the training image, but the sufficient statistics of the image are the mean values of the feature functions over the image. This is completely oblivious to variations across the image, and for this reason some small parts of the training image may actually have low probability. This is particularly true of textures composed of large randomly distributed textons (coherent texture elements like pebbles or grass blades), which greatly reduces the effective amount of training data, and thus is the most difficult to learn for all kinds of MGRF models. At one place in the image or over the whole image a particular angle or offset between textons may be more common than the average just by chance. Just like humans, sequential feature selection easily sees patterns in noise. We have already made the assumptions that the texture is homogeneous with the Markovian property and that the entire training image is a sample from the texture class, without foreign inclusions. We can thus expect that every piece of the training image of sufficient size (which we term the "scale" of the texture) is itself also a visually recognisable sample of the texture class, while those much smaller aren't, and thus the model should also recognise them. The scale will be approximately equal to once to twice the tesselation offset if the texture is regular, related to the size of any textons that the texture is composed of, and in general similar to the spatial extent of the Markov blanket. Indeed, this property holds for all of the training images used in our experiments. Hence it is unsatisfactory that only averages over the training image are considered, and we propose to switch away from the MLE as our objective.

Let $\{\mathbf{g}_i^\circ\}$ be a set of training images. In practice we simply split a single training image $\mathbf{g}^\circ$ into pieces of size 80×80 with overlap of 22 pixels. This is related to the maximum offset length between interacting pixels that we search for, 40 pixels. We change our optimisation problem to

$$\max_{p,\boldsymbol{\theta}} \min_{\mathbf{g}_i^\circ} \ell(\boldsymbol{\theta}|\mathbf{g}_i^\circ). \tag{4}$$

It is perhaps desirable to use a smoothed differentiable version of the objective using a 'soft-min' instead of min,

$$\max_{p,\boldsymbol{\theta}} \sum_i \ell(\boldsymbol{\theta}|\mathbf{g}_i^\circ) w(\mathbf{g}_i^\circ \mid p(\cdot|\boldsymbol{\theta})), \tag{5}$$

where w weights the pieces exponentially according to their proximity to the minimum (which is easy to compute):

$$\begin{aligned} w(\mathbf{g}_i^\circ|p) &:= \frac{\exp(-\alpha \ell(\boldsymbol{\theta}|\mathbf{g}_i^\circ))}{\sum_j \exp(-\alpha \ell(\boldsymbol{\theta}|\mathbf{g}_j^\circ))} \\ &= \frac{\exp(-\alpha \boldsymbol{\theta} \cdot \mathbf{h}(\mathbf{g}_i^\circ))}{\sum_j \exp(-\alpha \boldsymbol{\theta} \cdot \mathbf{h}(\mathbf{g}_j^\circ))} \end{aligned} \tag{6}$$

where α is a softness parameter.

This objective implies that when selecting features we replace the criterion

$$\arg\max_f d(\bar{\mathbf{h}}_f(\mathbf{g}_{\text{samp}}), \mathbf{h}_f(\mathbf{g}^\circ))$$

with one of

$$\arg\max_f \min_i d(\bar{\mathbf{h}}_f(\mathbf{g}_{\text{samp}}), \mathbf{h}_f(\mathbf{g}_i^\circ)) \tag{7}$$

or

$$\arg\max_f \sum_i w(\mathbf{g}_{\text{samp}}|p) d(\bar{\mathbf{h}}_f(\mathbf{g}_{\text{samp}}), \mathbf{h}_f(\mathbf{g}_i^\circ)) \tag{8}$$

to find features that are characteristic for all parts of the image.

4.2. Second order methods

As the data likelihood in the base model $p_i(\mathbf{g}^\circ)$ is fixed w.r.t. $\boldsymbol{\theta}$ and can be dropped, parameters of general exponential family distributions are learned in the same was as without a base model. The gradient vector $\nabla \ell(\boldsymbol{\theta}|\mathbf{g}^\circ)$ and the Hessian matrix $\mathbf{H}(\ell)(\boldsymbol{\theta}|\mathbf{g}^\circ)$ of the log-likelihood are easily derived as:

$$\nabla \ell(\boldsymbol{\theta}|\mathbf{g}^\circ) = -\mathbf{h}(\mathbf{g}^\circ) + E_{\boldsymbol{\theta}}\{\mathbf{h}(\mathbf{g})\} \tag{9}$$

$$\mathbf{H}(\ell)(\boldsymbol{\theta}|\mathbf{g}^\circ) = -\text{Cov}_{\boldsymbol{\theta}}\{\mathbf{h}(\mathbf{g})\} \tag{10}$$

where $E_{\boldsymbol{\theta}}\{\mathbf{h}(\mathbf{g})\}$ and $\text{Cov}_{\boldsymbol{\theta}}\{\mathbf{h}(\mathbf{g})\}$ denote respectively the expected value and covariance matrix of the image features $\mathbf{h}(\mathbf{g})$ w.r.t. the distribution $p(\mathbf{g}|\boldsymbol{\theta})$. As the covariance matrix is always non-positive definite, the log-likelihood is concave down in the space of parameter vectors, which is a common property of all exponential family distributions [56].

Computing the expectation in Eq. (9) exactly is in general intractable. Its approximation using a Markov chain Monte Carlo (MCMC) sampler such as the single-site Gibbs sampler can be highly unreliable, because the energy landscape often has deep local minima which are inescapable in reasonable timescales.

The Hessian may be reasonably and easily approximated as a diagonal matrix, so given samples from the model $p(\mathbf{g}|\boldsymbol{\theta}_j)$ we can find the 2nd order Taylor expansion about $\boldsymbol{\theta}_j$ as an approximation to the log-likelihood using Eq.s (9) and (10). A Newton step can be performed by inverting the Hessian, while a more numerically stable alternative is to find the maximum of the 2nd order likelihood approximation along the direction of the gradient. See [57] for details.

We use this second-order method to find an initial approximation to the parameters at each nesting iteration, to save a little time. As both gradient and Hessian are highly approximate and Newton's method takes dangerously large steps we take a single 2nd order step, and then switch to a stochastic 1st order method which follows only the approximate gradient, as detailed in the next section.

4.3. Controllable Simulated Annealing

Assuming that the time to sample from a MGRF is linear in the number of features, the nesting procedure runs in time quadratic in the number of nesting iterations. In order to keep running times reasonable, the number of nesting steps should be limited, and computation per nesting iteration minimised.

The most obvious method to learn the parameters of a MGRF is to perform stochastic gradient descent (SGD), following a noisy approximation to the gradient given by sampling from the model to approximate the expectation in the gradient (Eq. (9)). Sampling could be performed by starting from an initial image (e.g. of noise) and performing Gibbs sampler sweeps over the whole image until its total energy converges. However, this method is very slow. The MGRFs we learn are very dense, which makes sampling problematic; Gibbs sampling may mix very slowly, and easily gets stuck in local minima. Alternatively, a fixed number of sweeps can be performed to produce a sample, which is still useful for synthesising images but optimises the wrong objective.

Rather than repeatedly sampling from the models to approximate the gradient we used controllable simulated annealling (CSA) [58], simultaneously invented in [14], to produce approximate images matching the training statistics much faster than possible learning MLE parameters. An essentially identical procedure was also reinvented in [59] under the name 'persistent contrastive divergence' (PCD). The goal in PCD is parameter learning, and it is presently the most popular method for that for MGRFs. The difference is that PCD uses small step-sizes, and is started from the training data rather than a random image.

CSA alternates Gibbs sampling sweeps (visiting each pixel in the image $\mathbf{g}_{\text{samp}}$ once) and changes to parameters according to

$$\Delta\boldsymbol{\theta} = \boldsymbol{\lambda}_t \circ (\mathbf{h}(\mathbf{g}_{\text{samp}}) - \mathbf{h}(\mathbf{g}^{\circ}))$$

where $\circ$ denotes element-wise Hadamard product, $\mathbf{g}^{\circ}$ is the training image, $\boldsymbol{\lambda}_t$ is a vector of the current per-parameter step sizes and t is the iteration number. A Robbins-Monro (RM) sequence $\boldsymbol{\lambda}_t = \frac{15}{15+t}\mathbf{1}$ was used, where $\mathbf{1}$ is the all-ones vector of the same dimension as $\boldsymbol{\theta}$. CSA can more rapidly produce an image with statistics which usually match the training statistics fairly well, but also has a strong tendency to oscillate. For this reason we introduce a variant of CSA called accelerated controllable simulated annealing (ACSA), which uses an gain vector adaptive step size instead of RM steps, as described by Almeida et al. [60], with default parameters and initial step size vector $\boldsymbol{\lambda}_0 = \mathbf{1}$. This uses different learning rates for different parameters, dampening oscillations by reducing the learning rate for parameters which do so, while increasing the learning rate for those that seem to require it. Adaptive step sizes proved to give far faster convergence to the desired statistics than RM steps, and are more robust to initial step size.

Finding optimal parameters which produce good unguided synthesis results requires far more fine tuning to eliminate unwanted energy minima which are approached very slowly during MCMC. In theory, if optimal parameters are learned then MCMC sampling and CSA should both produce samples which have statistics equal to the constraints on average, and are indistinguishable by the model features. Hence CSA 'samples' can be substituted for real samples from the MLE model. Ideally if available a more efficient sampling algorithm than single-site Gibbs sampling should be used, but most popular samplers are not applicable here. More efficient algorithms have been developed for creating images matching desired statistics by making modifications attempting to directly move towards them [12, 13, 61].

4.4. Sampling in practice

Naturally, CSA will converge to the desired statistics faster if starting from parameters close to the MLE. At each nesting iteration we first found an approximation to the parameters as described in Section 4.2, and then ran CSA four times for 50 Gibbs sweeps each, starting from uniform noise images, to produce four samples of size 100×100. CSA results can be noisy and unreliable, not reaching the desired statistics, so it helps greatly to average over multiple runs. The parameters produced as a side effect of CSA are very approximate and may actually diverge from the MLE, but nonetheless are useful and carried forward to future iterations.

An important consideration when drawing samples from a MGRF texture model using an MCMC sampler is the choice of initial image. While the Markov chain converges asymptotically to the model distribution, this can take a number of steps exponential in the number of pixels due to ubiquitous deep local minima. However it is well known that MCMC converges slowly when highly correlated variables are updated independently, and in image models all nearby pixels are typically highly correlated to one another, making Gibbs sampling inefficient and prone to being trapped in local minima. The most common example of these are 'crystallisation' of regular patterns growing from two or more areas and meeting without aligning. Gibbs sampling of regular textures starting from noise often leads to crystallisation if the image size is several multiples of the maximum clique size.

There are a number of ways to avoid crystallisation. One is to add longer range interactions in order to more strongly enforce global structure. Otherwise, the process can be pushed towards the correct regularity by starting from an initial image with a template, in the form of dots, stripes, or any other unevenness of the desired tesselation. Such tesselation of near-regular textures can automatically be extracted by computing co-occurrence statistics for different offsets, in the form of a model-based interaction map (MBIM) [57]. Other alternatives are to slowly grow the image by extending the boundaries, starting from a single small 'seed' which is first allowed to converge [51], or starting from a random piece of the training image as a seed. The latter was used for these experiments due to simplicity. If a texture model is to be used for other tasks such as segmentation or classification better performance could be expected if appropriate learning, sampling and validation methods are used. For example for texture synthesis we care only about the local mimimum that is reached on sampling from the starting sample (a white noise image in this work), and disregard the

energy landscape outside of this basin, while for texture classification we want to ensure correct behaviour for distant images too.

For unguided texture synthesis after learning a model we used 300 sweeps of ACSA, growing from a seed, to produce images of size 180×180 plus trimmed margins of up to 65 pixels depending on feature sizes. Hence when inpainting this shortcut is not available: in practice CSA and ACSA are harmful for inpainting as generalisation (e.g. differing lighting) may be required. Instead we tuned the model parameters in the normal way, and then used Gibbs sampling to inpaint.

4.5. Binary patterns

In order to reduce the space of candidate potentials, we define families of potentials that are parameterised only by the shape of their supports, with unparameterised feature functions. The simplest such order k feature on an image of Q pixel grey-levels is the trivial grey level co-occurrence (GLC) feature

$$f_{\mathrm{GLC}}(x_1, \ldots, x_k) := x_1 Q^0 + \ldots + x_k Q^{k-1}$$

with Q^k bins. GLC features proved to have poor generalisation ability; while they performed well for texture synthesis, when used for inpainting tasks they remembered the original contrast and offset of the training image rather than matching the contrast of the surrounding inpainting frame. Previous researchers have found that histograms of the pairwise grey level difference (GLD) — defined for $k = 2$ as $f_{\mathrm{GLD}}(x_1, x_2) := x_2 - x_1$, with $2Q - 1$ bins — encodes the large majority of the information in a pairwise GLC histogram [16, 62], confirmed by our own comparisons GLC and GLD features for texture synthesis (see Section 5.1). Hence we used GLD instead of GLC features to capture pairwise interactions.

The high-order features investigated were binary patterns (BPs), generalising LBPs by using learned offsets from the central pixel. The BP feature function thresholds the grey levels of the pixels in each clique against a certain distinguished pixel (no longer necessarily in the centre): $f_{\mathrm{BP}}(x_0, \ldots, x_d) := \sum_{1 \le i \le d} 2^{i-1} [\![x_0 < x_i]\!]$. No interpolation between pixels was performed, nor were histogram bins combined as in uniform LBPs [28].

We write GLC^k, GLD^k and BP^k to indicate order k grey level co-occurrence, grey level difference and binary pattern features, respectively.

Binary equality features. Complementary to BPs, we considered 'binary equality' (BE) feature functions which test whether pixels are within an equality threshold c of each other (dependent on the grey level range) defined as

$$f_{\mathrm{BE}}(x_0, \ldots, x_d) := \sum_{1 \le i \le d} 2^{i-1} [\![|x_0 - x_i| \le c]\!].$$

We supposed these might be more suitable for describing flat regions of an image, where any amount of noise causes f_{BP} to produce evenly distributed random values. For synthesis purposes, these features must be used in heterogeneous models paired with other features that break the symmetry between an image and its inverse (GLD features fail to do this for images with both 180° rotational symmetry and symmetric histograms). In our experiments we found that if a texture contains flat regions then a number of BE features would be selected by the nesting procedure, otherwise they were hardly selected. We found them to have similar descriptive power to BP statistics, but had additional failure conditions due to the necessary symmetry breaking, so they were not used further.

4.6. Building up features

As a base model, we used a MGRF with a first order factor to model marginal statistics and the two nearest neighbour GLD features with offsets $((0,0),(1,0))$ and $((0,0),(0,1))$. As an exception, marginal potentials were not used for the inpainting experiment in Section 5.2, which improved the models' ability to match the contrast of the surrounding frame. However for unguided synthesis matching the original histogram is desirable.

All features were constrained to clique families with at most 40 pixels distance between two points. The initial candidate set (selector) C^1 was always comprised of all GLD^2 features within this circular window. Three GLD^2 features are added at a time to speed learning. Thereafter different possibilities for C^2 were considered as listed below.

- 'Combined' BP^5 features built out of the set of characteristic offsets $\{r_i\}$ occurring by the previously selected GLD^2 and BP^5 features. Each possible selection of two offsets r_1, r_2 and for each $i = 1, 2$ each choice of either using mirrored offsets $r_i, -r_i$ or halved offsets $r_i/2, -r_i/2$ provided the set of four-offset candidates. BP^5 features were also added two at a time.

- 'Conjoined' BP^9 features built by selecting (by exhaustive search) four BP^3 features with symmetric pairs of offsets $(r_i, -r_i)$ of maximal JSD between training and sample histograms and then combining them into a single clique family with shape $(r_1, -r_1, \ldots, r_4, -r_4)$. Examples can be seen in Figure 3.

- 'Jagged star' (jag-star) BP^k features (for $k = 9,13$) with $k-1$ surrounding pixels alternately and evenly spaced on two circles of radii d_0, d_1 centred on the central pixel, with offsets at

$$\begin{bmatrix} x_i \\ y_i \end{bmatrix} = d_{i \bmod 2} \begin{bmatrix} \cos(\frac{2\pi i}{k-1} + \phi) \\ \sin(\frac{2\pi i}{k-1} + \phi) \end{bmatrix},$$

  where $d_{i \bmod 2}$ is d_1 if i is odd and d_0 if i is even. These configurations include square as well as circular patterns. Considering a dense subset of rotations and radii from 1 to 20 pixels provided a fixed set of 2638 candidates. Examples can be seen in Figure 3.

- Linear filters. This is a fixed bank of Gabor wavelets and Laplacian of Gaussians identical to those used in FRAME [14] except that we increased the number of Gabor wavelet orientations from 6 to 10, and restricted sizes

of the filters to prevent excessively long running times. We employed Gabor wavelets with wavelengths only up to 6 pixels (filters of size 17×17) rather than 12. In total there were 64 filters.

5. Experiments

Generative image models have often been evaluated by measuring performance of image denoising and inpainting when the model is used as the prior. But this is an indirect and incomplete method of evaluation. Instead we mainly compare different classes of models by inspection of synthesis results. Experiments were conducted with a set of grey-scale digitised photographs of natural and approximately spatially homogeneous textures sourced from several popular databases. Textures were selected which were diverse, difficult to model, homogeneous and without periodicity or other features on a scale longer than 40 pixels (all images were kept at original scales, with the exception of the inpainting experiment in Section 5.2). The databases used were the popular Brodatz photoalbum [53]; the "NewTex" dataset of natural textures included with MeasTex [63], a framework for standardised texture discriminator evaluation; and the MIT VisTex database [54], and images collected by the NYU Laboratory for Computational Vision [64].

Textures were subjectively categorised into three classes according to structure: stochastic (those apparently described by only simple local interactions, e.g. sand, water), near-regular (regularly tiled but with random defects, e.g. weaves), and irregular (containing large scale elements with unpredictable placement or shapes, e.g. marbles).

Each image was quantised to $Q = 8$ grey levels. With some exceptions (e.g. when comparing to other results) the images were quantised using contrast-limited adaptive histogram equalization [65] with 16×16 tiles and a contrast clipping limit of 0.03. Adaptive histogram equalization mostly avoids the situation where most of the image is mapped to only one or two grey levels, and also lessens shadows and gradients, which hinder the recovery of long range interactions. The low Q value used speeds up Gibbs sampling.

Source code for all experiments is freely available from the website accompanying the paper, together with supplementary material including results for a large number of additional textures[1].

5.1. Synthesis

As a baseline we compare synthesis results to the texture analysis and synthesis algorithm by Portilla and Simoncelli [13] (downloadable from [64]), which in addition to being freely available, is today still one of the most successful attempts to model texture statistically, and has recently been extended to colour textures. Their approach uses iterated projections to attempt to produce images matching certain first- and second-order statistics of wavelet responses. On the other hand, all of the recent filter-based MRF texture models [9, 19–21, 66] have only been demonstrated on a few simple highly regular textures with short repetition lengths and texton scales. Unfortunately the lack of difficult synthesis examples or source code for these previously published approaches hinders comparison, although [21] achieved visually better results than us on the difficult D103/D104 textures (see Figure 9) and D76 by using three layers of hundreds of filters in a deep belief network. In theory the purpose of these layers was to stabilise learning by using tiled-convolutional rather than fully-convolutional filters, which for simpler models is not necessary [66].

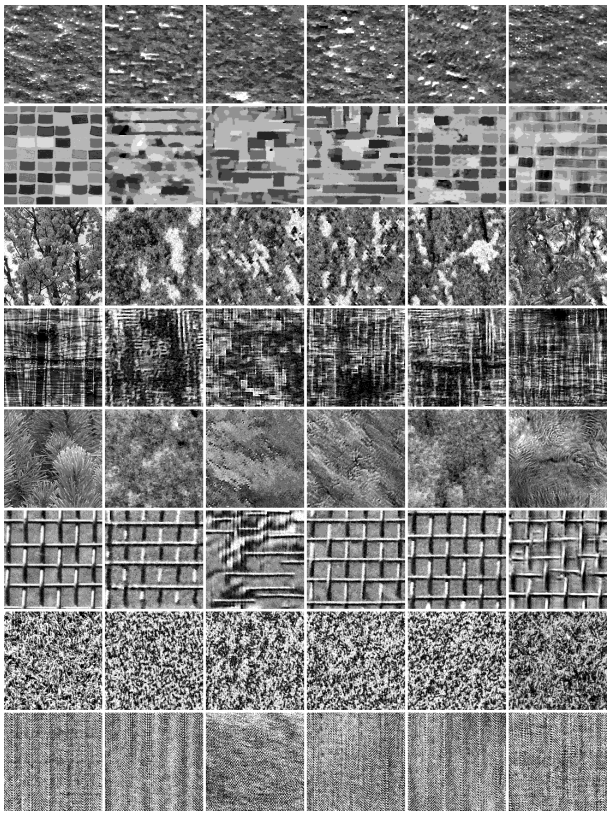

Figure 7: Comparison of synthesis results using FRAME linear filters to other models. Column 1: original textures; columns 2 to 5: models with GLD^2, linear filters, GLD^2 + linear filters, and GLD^2 + jag-star BP^{13} features, respectively.

Figures 4 and 5 show 180×180 synthesis results for 30 of the most diverse and interesting examples modelled using each of the alternative classes of BP features described in Section 4.6. Training images were 256×256 in size.

Synthesis results for the different models under comparison were evaluated by a panel of observers. For each of 20 textures (a subset of the 30 shown) six synthesised images from different classes of models was presented to each person along with the original training image, and they were tasked with placing the synthesis results in order from best to worst[2]. The observers were instructed to make rapid decisions rather than perform careful inspection.

Table 1 presents the average rankings of each of the model

[1] http://www.ivs.auckland.ac.nz/texture_modelling.

[2] This questionnaire is still available online at http://www.ivs.auckland.ac.nz/texture_modelling/.

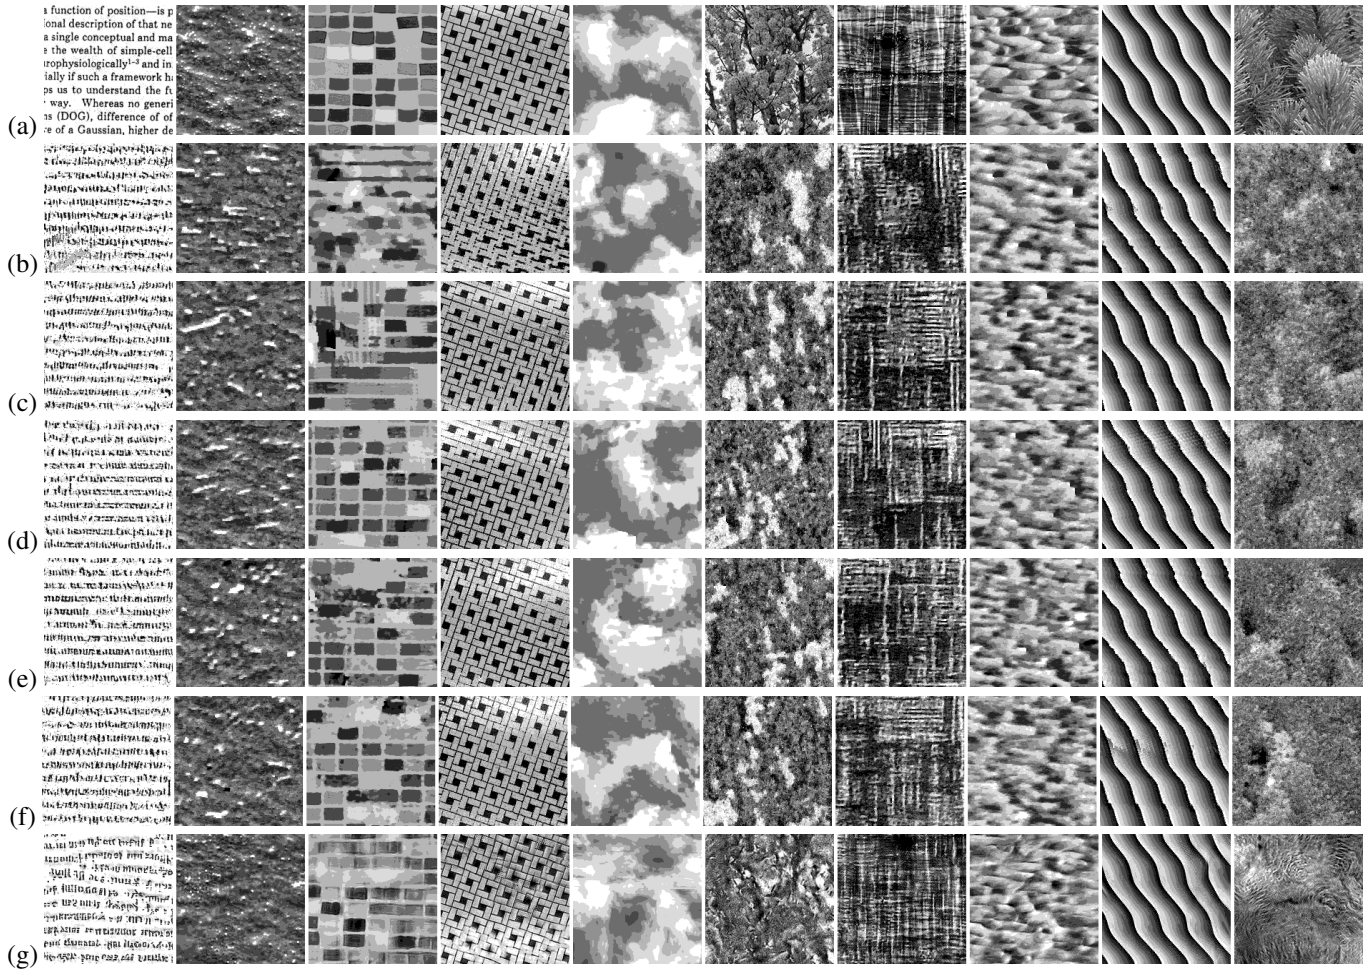

Figure 4: Synthesis results with models with differing kinds of potentials: (a) original textures; (b) second-order GLDs (GLD^2); (c) GLD^2+ 'combined' BP^5; (d) GLD^2+'conjoined' BP^9; (e) GLD^2+jag-star BP^9; (f) GLD^2+jag-star BP^{13}; and (g) Portilla and Simoncelli's results.

types against the other five options, summarised by texture category. Additionally Table 2 reports how frequently each class of model gave the best-ranked synthesis result. The BP^5 features outperforming the 9th order ones on stochastic textures may be an artifact of the earlier stopping rule, which usually excluded the higher order ones altogether on these textures.

From the figures and tables, it can be seen that the MGRFs perform well on near-regular textures, as tesselation is easily handled, and on stochastic textures, for which the apparent dependency of the actual texture is on a short scale and thus a natural fit for a MGRF. However, it is not easy to separate the performance of the different classes of models. On stochastic textures the method of Portilla and Simoncelli usually but not always dominates, apparently most significantly because it captures high frequency aspects and details (such as thin branches and speckles) that ours do not. It is often able to produce nice results for irregular textures too, however, it often violates fixed tesselation offsets and fails to represent complex textons. On those cases the simple and cheap GLD potentials can be very effective. In addition, Gibbs (or CSA) sampling from a MGRF adds unwanted high frequency noise. Complex irregular textures such as tiles in column 3 of Figure 4 are the most difficult for all methods, and in these cases hidden variables to distinguish different local contexts, as in gated MRFs, are probably necessary to visually replicate the texture.

In the tables the 'combined' 5^{th} order BPs are seen to do very well on irregular textures, despite being low order and only having 16 parameters per feature. This provides promising evidence that piecing together high order features out of low order ones is a workable technique. However a number of other schemes for building up to higher orders that we attempted, not documented here, provided bad results. This motivated the investigation of the jag-star features as simpler and less brittle alternatives. Previous independent experiments [51] also favoured the 'combined' features, providing more evidence for a real effect. No signficant difference between the 9^{th} and 13^{th} order jag-star features is visible, which is provides some evidence that the large number of parameter in the latter models — 4096 parameters per BP feature — is not unworkable. Results for star BP's (not shown) were close to those for jag-star, but generally inferior.

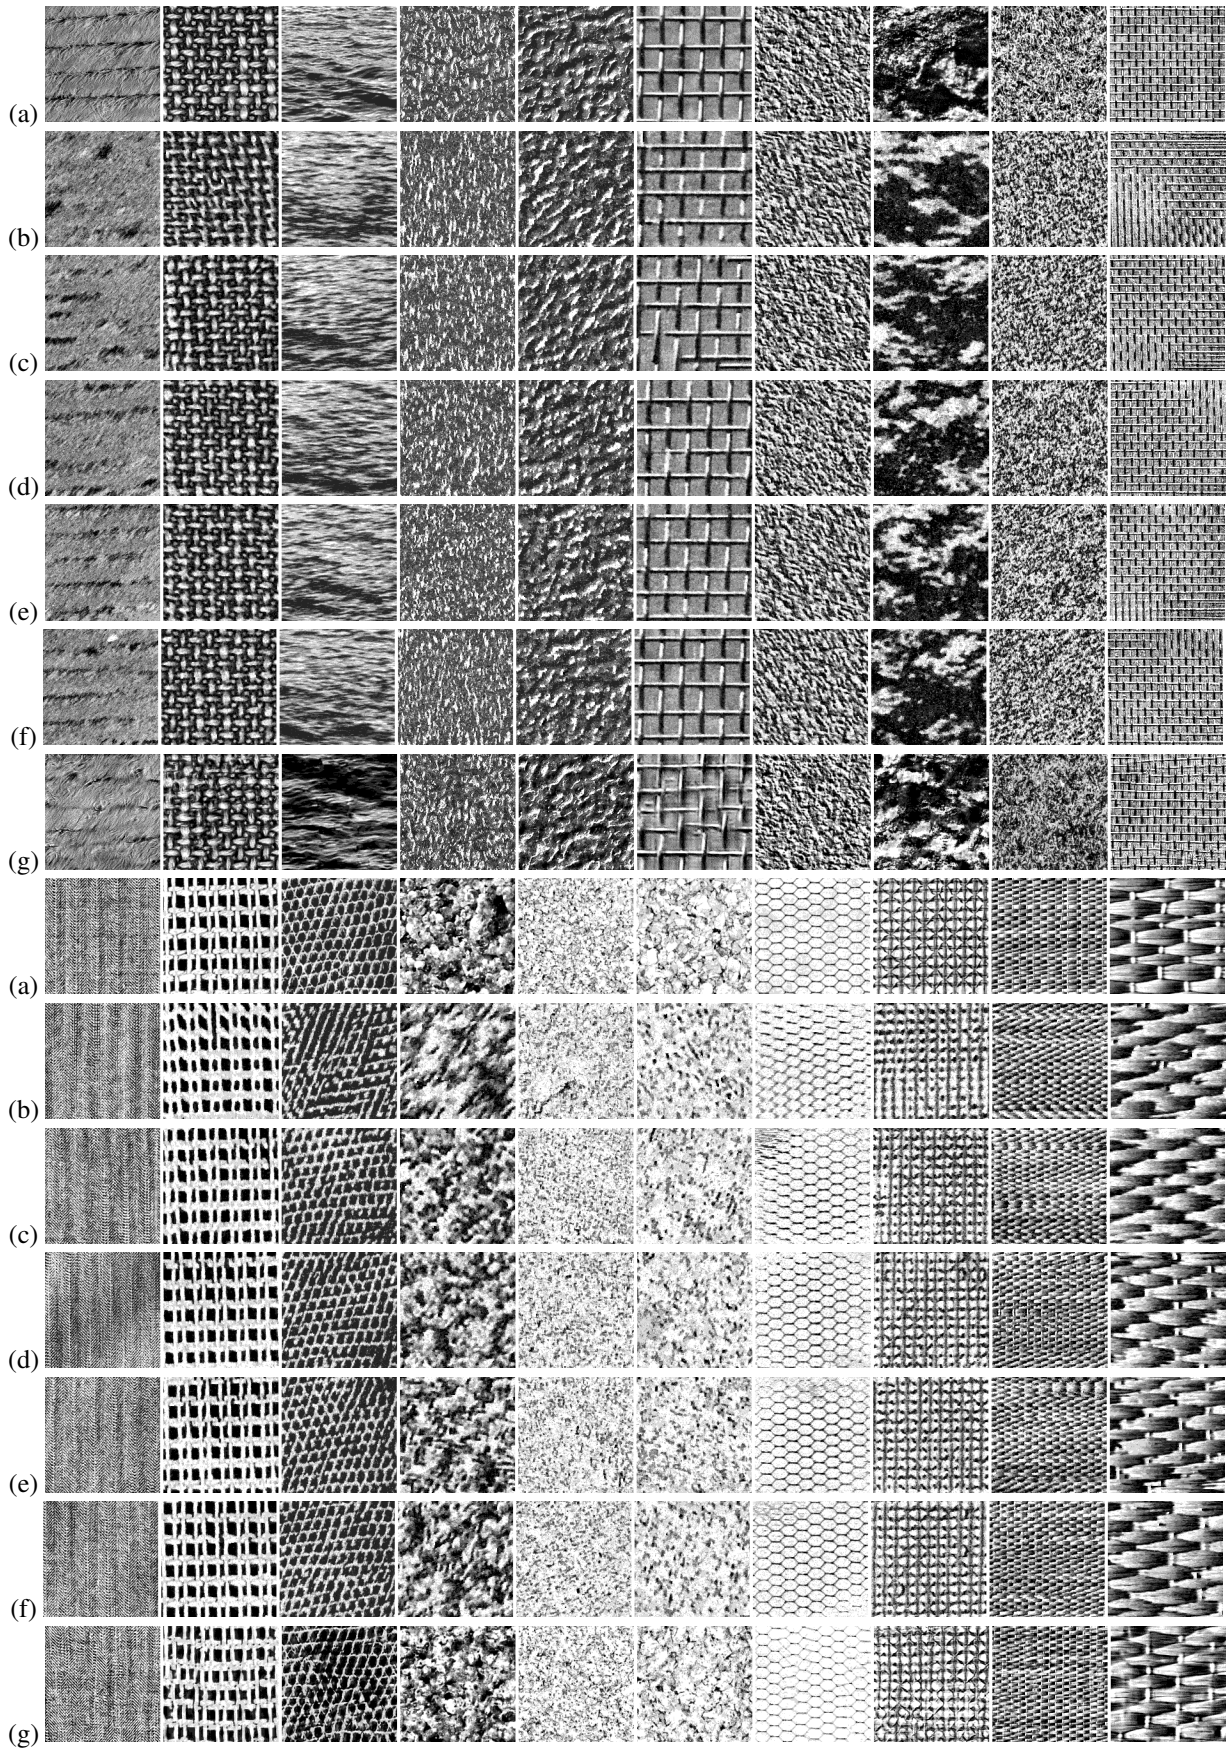

Figure 5: Synthesis results continued: see Figure 4.

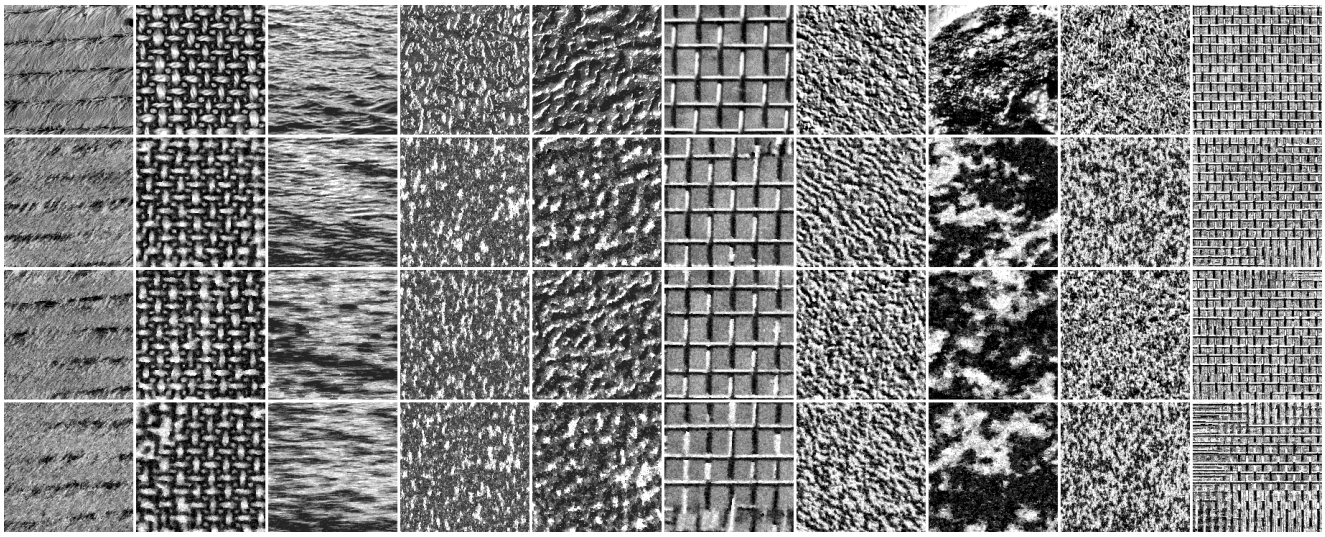

Figure 6: Variants on jag-star BP models. Rows are: *First*: training images. *Second*: GLD^2 + jag-star BP^9 models. *Third*: GLC^2 + jag-star BP^9 models. *Fourth*: jag-star BP^9 models without any pairwise potentials (except for initial adjacent ones).

Table 1: Averages and standard deviations of human rankings of texture synthesis results by texture category: 1 was best and 6 worst. All models except Portilla & Simoncelli [13] include GLD potentials.

	Stochastic	Near-regular	Irregular	Total
Lin. filters	3.0±1.3	4.1±1.2	4.0±1.4	3.8±1.3
Comb. BP^5	4.1±1.2	2.0±1.0	3.6±1.3	3.2±1.2
Conj. BP^9	4.4±1.3	3.4±1.2	3.0±1.2	3.5±1.2
JStar BP^9	4.0±1.5	3.8±1.2	3.7±1.2	3.8±1.3
JStar BP^{13}	4.0±1.3	3.3±1.0	3.1±1.1	3.4±1.1
P&S [13]	1.5±1.0	4.3±1.3	3.7±1.4	3.3±1.3

Table 2: Fraction of the time that an image synthesised by each class of texture model was ranked best by a judge out of the 6 possibilities, split by texture category. All models except Portilla & Simoncelli [13] include GLD potentials.

	Stochastic	Near-regular	Irregular	Total
Lin. filters	6%	9%	11%	9%
Comb. BP^5	2%	47%	11%	21%
Conj. BP^9	0%	12%	22%	13%
JStar BP^9	6%	7%	5%	6%
JStar BP^{13}	2%	12%	26%	15%
P&S [13]	82%	9%	23%	33%

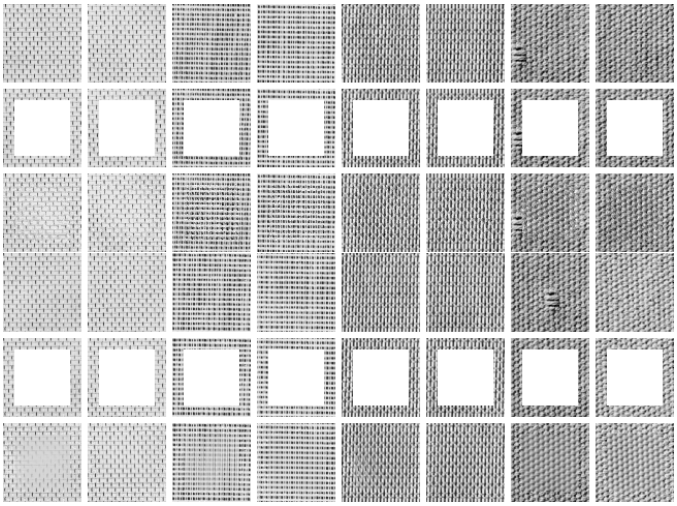

Figure 8: Inpainting examples for a 9th-order jag-star model of D6, D21, D53 and D77. The left hand column of each pair shows the worst inpainting result over 20 repetitions, the right hand column the best result. The top block of results is for simple inpainting using Gibbs sampling, the bottom block shows a smoothed result by averaging 50 consecutive Gibbs iterations. The smoothed images have much higher MSSIM scores.

5.1.1. GLC vs. GLD features

We also compared the effect of replacing pairwise GLD features with GLC features or removing them entirely, as shown in Figure 6. Generally the result of excluding pairwise potentials is not very great (although it is very large when linear filters are used instead, Section 5.1.2), or even undetectable, as the higher order features can seemingly usually capture the same statistics, which is against expectations. Naturally long range GLD features are very helpful for regular textures. The difference between GLD and GLC features is even less, though over the whole texture set use of GLC models as a base seem to be slightly more powerful.

5.1.2. Linear vs. nonlinear filters

Using BPs for texture modelling was compared with more traditional linear filters, too. We consider both models which are close to those used in [14] (with both nesting procedure and filter bank being similar), and models also including GLD^2 features. In either case the initial base model was the same as used in all our other models.

We can see from Figure 7 that small linear filters alone of course can not reconstruct textures whenever textural elements are larger than the filter wavelength, although larger objects can

still be partially captured, such as the tiles in row 2, column 3. Augmenting the filters with long range GLD potentials prevents this failure case, while the linear filters perform much better than pairwise features alone. Unexpectedly with the addition of these long range potentials the synthesised images do not normally appear distinct from those composed of BP features, with very similar results across all texture types. For example the speckles in the first row texture are disappointingly not captured, indicating a deficiency in the fixed filterbank. However in row 5 (pine trees) we do see some ripples in the images clearly captured by wavelets, including Portilla & Simoncelli's analysis, but not by BP features. Additionally from Table 1 it appears that the FRAME linear filters outperform BPs on stochastic textures, as might be expected, although far weaker than the comprehensive set of wavelet statistics used in [13]; but on the other hand are outperformed by BPs on textures with more structure.

Despite restricting the filter sizes to 17×17, model learning times with the filters were still 30 minutes on average (our implementation in C++ is single-threaded), several times that of the BP models.

5.2. *Inpainting*

Additionally, we compared against results from three recent leading hierarchical MGRF texture models [19, 21, 66] which attempted to quantitatively evaluate MGRF texture models through a texture inpainting task. (Heess et al. [9] used a slightly different variant, so most of their results are incomparable.) By focusing on four highly regular textures with small textural feature sizes, it is reasonable to compare the synthesized pixels in the inpainted region to the original ones since they should be nearly aligned. This is not applicable to other texture types, so is a very narrow test.

In the task a 76×76 pixel piece of Brodatz textures D6, D21, D53, or D77 is provided as an inpainting frame with a 54×54 pixel hole cut out of the centre. An algorithm fills the hole in a way consistent with the frame, and the difference against the ground truth is measured using the mean structural similarity index (MSSIM) [67]. MSSIM has been found experimentally to be a good measure of perceptual distance between images. D6 and D53 were bilinearly scaled to 50% and D21 and D77 to 75%; this allowed the use of small filters by the earlier papers. The top half of each image was used for training, and the bottom half for testing. Since the learning procedure is stochastic, the inpainting is repeated (200 times in our case) with random inpainting frames from test region and the average MSSIM score reported.

Examples of the inpainted frames are shown in Figure 8, and examples of texture *synthesis* using the models used for inpainting is shown in Figure 9, comparing to previously published synthesis results. Initially, the MSSIM scores across all models we learnt were very poor, but Figure 9 reveals the reason for the bad quantitative results. All the best previously published models produce very smooth synthesis and also inpainting (not shown) results, as if maximising the per-pixel posterior marginals rather than providing a sample from the texture distribution. This is especially true for the published works which achieve the highest MSSIM scores [66]. On the other hand our models produce typically more realistic-appearing inpainting and synthesis results, with random deviations.

By averaging over 50 consecutive Gibbs MCMC samples we now achieve similar scores to previously published algorithms, although visual quality is greatly decreased by doing so. This is demonstrated in Figure 8. The MSSIM scores are presented in Table 3, alongside best previously published results and the Efros and Leung texture synthesis algorithm [42] as a baseline. For this experiment we used $Q = 16$ rather than $Q = 8$ grey levels. The comparison of ground truth with 256 grey levels against models using only 16 levels gives a slight penalty against the previous works. For some unknown reason all models attempted frequently failed to properly inpaint D6, resulting in low scores for this texture. The full resolution version of D6 is also unusually difficult to synthesize for all models.

6. Conclusion

While Markov random fields have long been applied to image modelling, there has been little prior investigation in this context into practical learning of interaction structure, with until recently [33] apparently only 2nd to 4th order MGRFs previously learned. The particular model nesting approach applied in this paper has the same theoretical foundation on maximum entropy distributions as previous sequential structure learning approaches, but emphasizes approximation and speed by using CSA. Aside from this, the strength of model nesting is in creating heterogeneous models, and allowing composition of feature functions. Mixing dissimilar potentials in a MGRF typically shows immediate improvements.

Furthermore, for nearly two decades MGRF texture and image models based on square linear filters have been almost exclusively studied due to their superior performance over the primitive pairwise models that preceded them. However, we have shown that other types of high-order features, such as the generic binary pattern features introduced in this paper, are viable alternatives to linear filters. These BP features have several advantages over filters. They are offset and nearly contrast invariant, and are still powerful even when of low order. The presented synthesis and inpainting experiments showed that even BP feature of only 5^{th} order are sufficient to capture many important visual aspects of textures and are capable of outperforming the previously best approaches, thanks to building features up out of characteristic low-order cliques. These relatively low orders reduce computation costs, while no filter coefficients need to be selected or learned.

The purpose of our synthesis experiments is not to challenge the mainstream highly effective and efficient texture synthesis algorithms but to validate our probabilistic models. There are however still many textures which our models have failed to capture, especially complex irregular ones, as well as losing fine details in general. In order to tackle these in future we plan to extend our models by including additional layers of feature selectors for diverse types of statistics, particularly co-occurrences between the outputs of feature functions at previous levels, as has been used in several texture models which

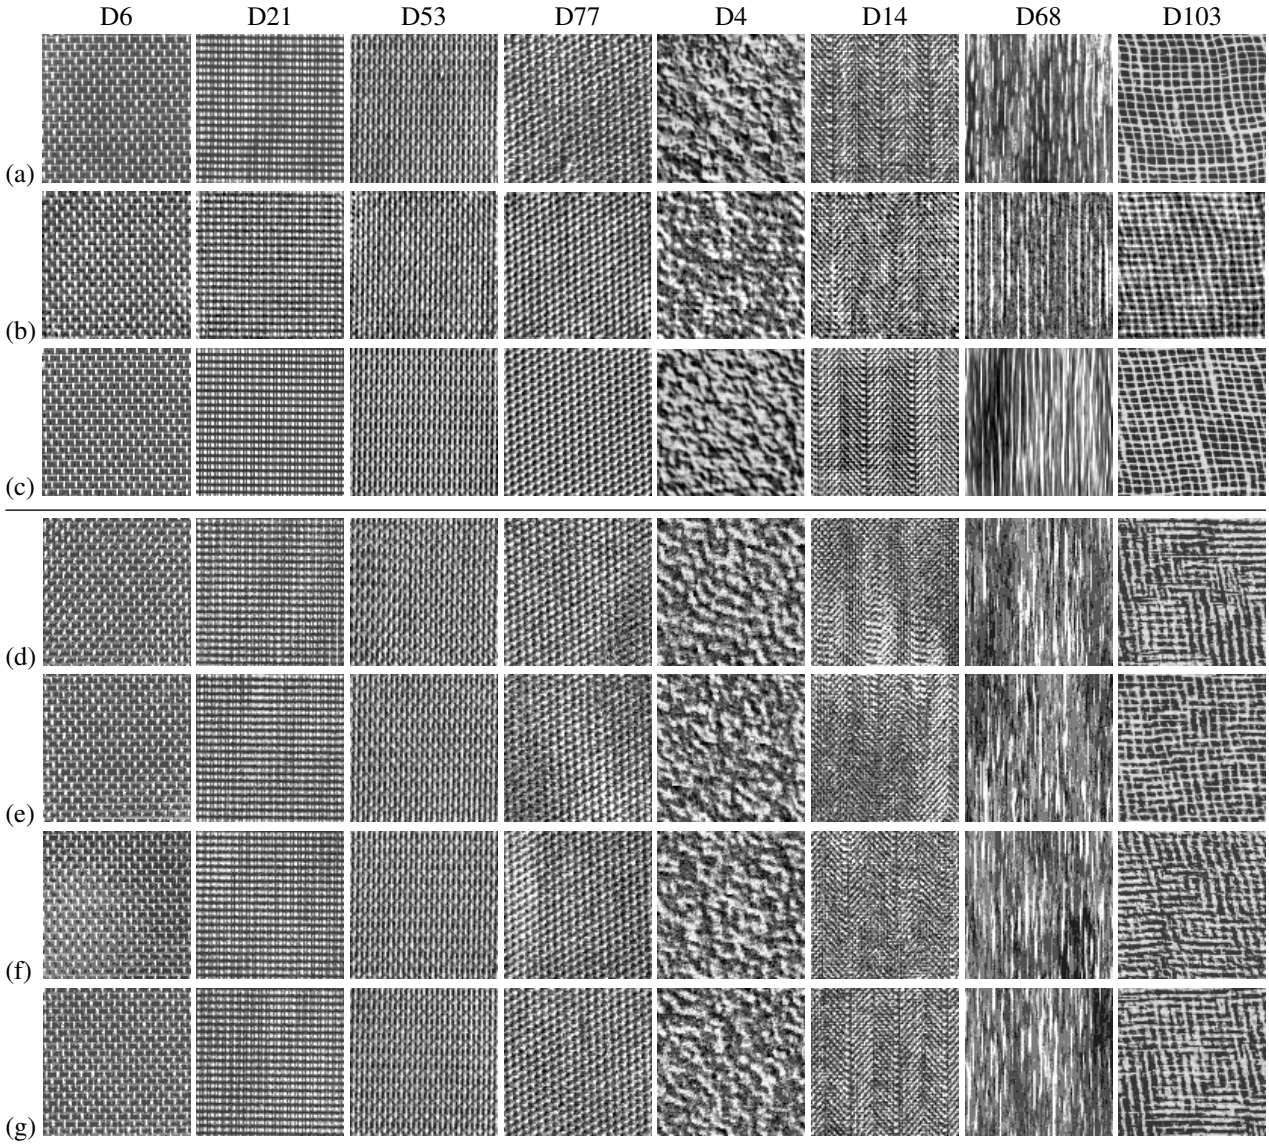

Figure 9: Comparison of our synthesis results against previously published works: (a)the eight original 98×98 Brodatz textures; (b) results of the Multi-Tm algorithm [19]; (c) results of the 2-layer DBM [21]; (d) – (g) our results with GLD^2, GLD^2 + combined BP^5, GLD^2 + conjoin BP^9, and GLD^2 + jagstar BP^{13}. The original images are scaled and with grey levels reversed as in the previous works. Each image has been individually renormalised in order to erase contrast differences in the original textures.

	D6	D21	D53	D77
Efros & Leung [42]	0.85 ± 0.03	0.86 ± 0.03	0.86 ± 0.06	0.60 ± 0.08
TmPoT [19]	0.86 ± 0.02	0.87 ± 0.01	0.86 ± 0.02	0.77 ± 0.03
TssRBM [21]	0.89 ± 0.02	0.91 ± 0.01	0.92 ± 0.02	0.76 ± 0.03
2-layer DBN [21]	0.89 ± 0.03	0.91 ± 0.02	0.92 ± 0.03	0.77 ± 0.02
cGRBMs [66]	0.91 ± 0.02	0.93 ± 0.01	0.93 ± 0.01	0.78 ± 0.03
GLD^2	0.66 ± 0.15	0.90 ± 0.04	0.86 ± 0.08	0.78 ± 0.03
Linear filters	0.77 ± 0.02	0.49 ± 0.13	0.89 ± 0.04	0.58 ± 0.06
Combined BP^5	0.80 ± 0.08	0.87 ± 0.06	0.92 ± 0.03	0.81 ± 0.02
Conjoined BP^9	0.77 ± 0.10	0.86 ± 0.06	0.90 ± 0.03	0.80 ± 0.03
Jag-star BP^9	0.84 ± 0.08	0.86 ± 0.07	0.91 ± 0.03	0.80 ± 0.03
Jag-star BP^{13}	0.73 ± 0.11	0.90 ± 0.05	0.91 ± 0.03	0.79 ± 0.03

Table 3: Results of inpainting evaluation; MSSIM scores against ground truth (*mean* ± *standard deviation*). All our models (in the bottom half) include pairwise GLD potentials in addition to the higher order ones. Results for the Efros & Leung algorithm are taken from [19] and used 19×19 patch sizes.

used filters as features. A stopping rule for the nesting procedure which is robust to the highly variable quality of CSA samples would also be important if more feature sets are used, to keep the number of nesting iterations low. In future work we also hope to improve the generalisation ability of the models, such as with invariance to small deformations of the image. This is closely related to the selection of complex composite features which can describe more powerful abstractions.

References

[1] A. M. Ali, A. A. Farag, G. L. Gimel'farb, Optimizing binary MRFs with higher order cliques, in: D. Forsyth, P. Torr, A. Zisserman (Eds.), Proc. European Conf. on Computer Vision (ECCV 2008), Vol. 5304 of Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2008, pp. 98–111.
URL http://dx.doi.org/10.1007/978-3-540-88690-7_8

[2] N. Komodakis, N. Paragios, Beyond pairwise energies: Efficient optimization for higher-order MRFs, in: IEEE Conf. on Computer Vision and Pattern Recognition (CVPR'09), 2009, pp. 2985–2992. doi:10.1109/CVPR.2009.5206846.

[3] C. Wang, N. Komodakis, N. Paragios, Markov random field modeling, inference & learning in computer vision & image understanding: A survey, Comput. Vis. Image Underst. 117 (11) (2013) 1610–1627. doi:10.1016/j.cviu.2013.07.004.
URL http://dx.doi.org/10.1016/j.cviu.2013.07.004

[4] V. Kolmogorov, R. Zabin, What energy functions can be minimized via graph cuts?, Pattern Analysis and Machine Intelligence, IEEE Transactions on 26 (2) (2004) 147–159. doi:10.1109/TPAMI.2004.1262177.

[5] G. R. Cross, A. K. Jain, Markov random field texture models, IEEE Trans. on Pattern Analysis and Machine Intelligence 5 (1) (1983) 25–39. doi:10.1109/TPAMI.1983.4767341.

[6] S. Geman, D. Geman, Stochastic relaxation, gibbs distributions, and the bayesian restoration of images, IEEE Transactions on Pattern Analysis and Machine Intelligence 6 (1984) 721–741.

[7] A. Gagalowicz, S. D. Ma, Sequential synthesis of natural textures, Computer Vision, Graphics, and Image Processing 30 (3) (1985) 289 – 315. doi:http://dx.doi.org/10.1016/0734-189X(85)90162-8.
URL http://www.sciencedirect.com/science/article/pii/0734189X85901628

[8] G. Gimel'farb, Quantitative description of spatially homogeneous textures by characteristic grey level co-occurrences, Australian Journal of Intelligent Information Processing Systems 6 (2000) 46–53.

[9] N. Heess, C. K. I. Williams, G. E. Hinton, Learning generative texture models with extended Fields-of-Experts, in: Proc. British Machine Vision Conf. (BMVC'09), 2009, pp. 1–11. doi:10.5244/C.23.115.

[10] S. C. Zhu, Y. Wu, D. Mumford, Minimax entropy principle and its application to texture modeling, Neural Computation 9 (8) (1997) 1627–1660.

[11] S. Roth, M. J. Black, Fields of Experts, Int. J. of Computer Vision 82 (2) (2009) 205–229. doi:10.1007/s11263-008-0197-6.

[12] D. J. Heeger, J. R. Bergen, Pyramid-based texture analysis/synthesis, in: Proc. 22nd Conf. on Computer Graphics and Interactive Techniques, SIGGRAPH '95, ACM, New York, 1995, pp. 229–238. doi:10.1145/218380.218446.
URL http://doi.acm.org/10.1145/218380.218446

[13] J. Portilla, E. P. Simoncelli, A parametric texture model based on joint statistics of complex wavelet coefficients, Int. J. Computer Vision 40 (1) (2000) 49–70. doi:10.1023/A:1026553619983.
URL http://dx.doi.org/10.1023/A:1026553619983

[14] S. C. Zhu, Y. Wu, D. Mumford, Filters, random fields and maximum entropy (FRAME): Towards a unified theory for texture modeling, Int. J. of Computer Vision 27 (2) (1998) 107–126.

[15] K. Valkealahti, E. Oja, Reduced multidimensional co-occurrence histograms in texture classification, Pattern Analysis and Machine Intelligence, IEEE Transactions on 20 (1) (1998) 90–94. doi:10.1109/34.655653.

[16] T. Ojala, K. Valkealahti, E. Oja, M. Pietikäinen, Texture discrimination with multidimensional distributions of signed gray-level differences, Pattern Recognition 34 (3) (2001) 727 – 739. doi:http://dx.doi.org/10.1016/S0031-3203(00)00010-8.

[17] B. Julesz, Visual pattern discrimination, Information Theory, IRE Trans. on 8 (2) (1962) 84–92. doi:10.1109/TIT.1962.1057698.

[18] S. Della Pietra, V. Della Pietra, J. Lafferty, Inducing features of random fields, IEEE Trans. on Pattern Analysis and Machine Intelligence 19 (4) (1997) 380–393. doi:10.1109/34.588021.

[19] J. J. Kivinen, C. Williams, Multiple texture Boltzmann machines, in: N. D. Lawrence, M. A. Girolami (Eds.), Proc. 15th Int. Conf. on Artificial Intelligence and Statistics, 2012, pp. 638–646 vol.2.
URL http://jmlr.csail.mit.edu/proceedings/papers/v22/kivinen12/kivinen12.pdf

[20] M. Ranzato, V. Mnih, G. E. Hinton, Generating more realistic images using gated MRFs, in: Advances in Neural Information Processing Systems 23 (NIPS'10), Curran Associates, 2010, pp. 2002–2010.
URL http://media.nips.cc/nipsbooks/nipspapers/paper_files/nips23/NIPS2010_1163.pdf

[21] H. Luo, P. L. Carrier, A. Courville, Y. Bengio, Texture modeling with convolutional spike-and-slab RBMs and deep extensions, J. of Machine Learning Research 31 (2013) 415–423.

[22] A. Hyvärinen, E. Oja, Independent component analysis: algorithms and applications, Neural networks 13 (4) (2000) 411–430.

[23] Y. W. Teh, M. Welling, S. Osindero, G. E. Hinton, Energy-based models for sparse overcomplete representations, J. of Machine Learning Research 4 (2003) 1235–1260.

[24] G. Sharma, S. Hussain, F. Jurie, Local higher-order statistics (LHS) for texture categorization and facial analysis, in: A. Fitzgibbon, S. Lazebnik,

P. Perona, Y. Sato, C. Schmid (Eds.), Computer Vision – ECCV 2012, Vol. 7578 of Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2012, pp. 1–12.

[25] K. Sivakumar, J. Goutsias, Morphologically constrained grfs: applications to texture synthesis and analysis, Pattern Analysis and Machine Intelligence, IEEE Transactions on 21 (2) (1999) 99–113. doi:10.1109/34.748817.

[26] W.-H. Liao, T.-J. Young, Texture classification using uniform extended local ternary patterns, in: Multimedia (ISM), 2010 IEEE Int. Symposium on, 2010, pp. 191–195. doi:10.1109/ISM.2010.35.

[27] T. Ojala, M. Pietikäinen, D. Harwood, A comparative study of texture measures with classification based on featured distributions, Pattern Recognition 29 (1) (1996) 51 – 59. doi:http://dx.doi.org/10.1016/0031-3203(95)00067-4.

[28] T. Ojala, M. Pietikäinen, T. Maenpaa, Multiresolution gray-scale and rotation invariant texture classification with local binary patterns, IEEE Trans. on Pattern Analysis and Machine Intelligence 24 (7) (2002) 971–987. doi:10.1109/TPAMI.2002.1017623.

[29] R. Nosaka, C. Suryanto, K. Fukui, Rotation invariant co-occurrence among adjacent lbps, in: J.-I. Park, J. Kim (Eds.), Computer Vision - ACCV 2012 Workshops, Vol. 7728 of Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2013, pp. 15–25.

[30] X. Tan, B. Triggs, Enhanced local texture feature sets for face recognition under difficult lighting conditions, in: S. K. Zhou, W. Zhao, X. Tang, S. Gong (Eds.), Analysis and Modeling of Faces and Gestures, Vol. 4778 of Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2007, pp. 168–182.

[31] Y. Zhai, D. Neuhoff, T. Pappas, Local radius index - a new texture similarity feature, in: IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP), 2013, pp. 1434–1438. doi:10.1109/ICASSP.2013.6637888.

[32] N. Liu, G. Gimel'farb, P. Delmas, Y. H. Chan, Texture modelling with generic translation- and contrast/offset-invariant 2nd-4th order MGRFs, in: Proc. 28th Int. Conf. on Image and Vision Computing New Zealand (IVCNZ 2013), 2013, pp. 370–375.

[33] N. Liu, G. Gimel'farb, P. Delmas, High-order mgrf models for contrast/offset invariant texture retrieval, in: Proceedings of the 29th International Conference on Image and Vision Computing New Zealand, IVCNZ '14, ACM, New York, NY, USA, 2014, pp. 96–101. doi:10.1145/2683405.2683414.

[34] A. Zalesny, Analysis and synthesis of textures with pairwise signal interactions, Tech. Rep. KUL/ESAT/PSI/9902, Katholieke Universiteit Leuven (1999).

[35] S. Perkins, K. Lacker, J. Theiler, I. Guyon, A. Elisseeff, Grafting: Fast, incremental feature selection by gradient descent in function space, J. of Machine Learning Research 3 (2003) 1333–1356.

[36] J. Zhu, N. Lao, E. P. Xing, Grafting-light: fast, incremental feature selection and structure learning of Markov random fields, in: Proc. 16th ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining, ACM, 2010, pp. 303–312.

[37] M. Schmidt, K. Murphy, Convex structure learning in log-linear models: Beyond pairwise potentials, in: Proc. Int. Workshop on Artificial Intelligence and Statistics, 2010, pp. 709–716.

[38] A. McCallum, Efficiently inducing features of conditional random fields, in: Proc. 19th Conf. on Uncertainty in Artificial Intelligence (UAI'03), Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2003, pp. 403–410.
URL http://dl.acm.org/citation.cfm?id=2100584.2100633

[39] L. Mihalkova, R. J. Mooney, Bottom-up learning of Markov logic network structure, in: Proc. 24th Int. Conf. on Machine Learning, ICML '07, ACM, New York, 2007, pp. 625–632. doi:10.1145/1273496.1273575.
URL http://doi.acm.org/10.1145/1273496.1273575

[40] S. Lee, V. Ganapathi, D. Koller, Efficient structure learning of Markov networks using L1-regularization, in: Advances in Neural Information Processing Systems 19 (NIPS'06), MIT Press, 2007, pp. 817–824.

[41] Y. Chen, M. Welling, Bayesian structure learning for Markov random fields with a spike and slab prior, in: Proc. 28th Conf. on Uncertainty in Artificial Intelligence, AUAI Press, 2012, pp. 174–184.

[42] A. A. Efros, T. K. Leung, Texture synthesis by non-parametric sampling, in: Proc. Seventh IEEE Int. Conf. Computer Vision, 1999, pp. 1033–1038 vol.2. doi:10.1109/ICCV.1999.790383.

[43] L.-Y. Wei, M. Levoy, Fast texture synthesis using tree-structured vector quantization, in: Proc. 27th Conf. on Computer Graphics and Interactive Techniques, SIGGRAPH '00, ACM Press/Addison-Wesley Publishing, New York, 2000, pp. 479–488. doi:10.1145/344779.345009.
URL http://dx.doi.org/10.1145/344779.345009

[44] A. A. Efros, W. T. Freeman, Image quilting for texture synthesis and transfer, in: Proc. 28th Conf. on Computer Graphics and Interactive Techniques, SIGGRAPH '01, ACM, New York, 2001, pp. 341–346. doi:10.1145/383259.383296.
URL http://doi.acm.org/10.1145/383259.383296

[45] V. Kwatra, A. Schödl, I. Essa, G. Turk, A. Bobick, Graphcut textures: Image and video synthesis using graph cuts, ACM Trans. on Graphics, SIGGRAPH 2003 22 (3) (2003) 277–286.

[46] K. Popat, R. W. Picard, Novel cluster-based probability model for texture synthesis, classification, and compression, in: Visual Communications '93, Int. Society for Optics and Photonics, 1993, pp. 756–768.

[47] R. Paget, I. Longstaff, Texture synthesis via a noncausal nonparametric multiscale Markov random field, IEEE Trans. on Image Processing 7 (6) (1998) 925–931. doi:10.1109/83.679446.

[48] G. Peyré, Sparse modeling of textures, J. of Mathematical Imaging and Vision 34 (1) (2009) 17–31. doi:10.1007/s10851-008-0120-3.
URL http://dx.doi.org/10.1007/s10851-008-0120-3

[49] E. T. Jaynes, Information theory and statistical mechanics, Phys. Rev. 106 (1957) 620–630. doi:10.1103/PhysRev.106.620.

[50] M. Dudík, S. J. Phillips, R. E. Schapire, Maximum entropy density estimation with generalized regularization and an application to species distribution modeling, J. Mach. Learn. Res. 8 (2007) 1217–1260.
URL http://dl.acm.org/citation.cfm?id=1314498.1314540

[51] R. Versteegen, G. Gimel'farb, P. Riddle, Learning high-order generative texture models, in: Proc. 29th Int. Conf. on Image and Vision Computing New Zealand (IVCNZ 2014), 2014, pp. 90–95.

[52] J. Lin, Divergence measures based on the Shannon entropy, IEEE Trans. on Inf. Theory 37 (1) (1991) 145–151. doi:10.1109/18.61115.

[53] P. Brodatz, Textures: a Photographic Album for Artists and Designers, Dover, New York, 1966.
URL http://www.ux.uis.no/~tranden/brodatz.html

[54] Vision Texture, http://vismod.media.mit.edu/vismod/imagery/VisionTexture/, accessed: 2013-09-02.

[55] J. Žujović, Perceptual texture similarity metrics, Ph.D. thesis, Northwestern University (2011).

[56] O. Barndorff-Nielsen, Exponential families, Wiley Online Library, 1982.

[57] G. Gimel'farb, D. Zhou, Texture analysis by accurate identification of a generic Markov–Gibbs model, in: H. Bunke, A. Kandel, M. Last (Eds.), Applied Pattern Recognition, Vol. 91, Springer Berlin Heidelberg, 2008, pp. 221–245.

[58] G. Gimel'farb, Image Textures and Gibbs Random Fields, Kluwer Academic Publishers, Dordrecht, 1999.

[59] T. Tieleman, Training restricted Boltzmann machines using approximations to the likelihood gradient, in: Proc. 25th Int. Conf. on Machine Learning (ICML'08), ACM, 2008, pp. 1064–1071.

[60] L. B. Almeida, T. Langlois, J. D. Amaral, A. Plakhov, Parameter adaptation in stochastic optimization, in: D. Saad (Ed.), On-line Learning in Neural Networks, Cambridge University Press, New York, 1998, pp. 111–134.
URL http://dl.acm.org/citation.cfm?id=304710.304730

[61] S.-C. Zhu, X. Liu, Y. Wu, Exploring texture ensembles by efficient Markov chain Monte Carlo – toward a "trichromacy" theory of texture, IEEE Trans. on Pattern Analysis and Machine Intelligence 22 (6) (2000) 554–569. doi:10.1109/34.862195.

[62] M. Unser, Sum and difference histograms for texture classification, IEEE Trans. Pattern Anal. Mach. Intell. 8 (1) (1986) 118–125. doi:10.1109/TPAMI.1986.4767760.
URL http://dx.doi.org/10.1109/TPAMI.1986.4767760

[63] Meastex image texture database and test suite, http://www.texturesynthesis.com/meastex/meastex.html, accessed: 2013-09-02.

[64] J. Portilla, E. P. Simoncelli, Representation and synthesis of visual texture, http://www.cns.nyu.edu/~lcv/texture/, accessed: 2014-10-10.

[65] K. Zuiderveld, Contrast limited adaptive histogram equalization, in:

Graphic Gems IV, Academic Press Professional, 1994, pp. 474–485.

[66] Q. Gao, S. Roth, Texture synthesis: From convolutional RBMs to efficient deterministic algorithms, in: P. Fränti, G. Brown, M. Loog, F. Escolano, M. Pelillo (Eds.), Structural, Syntactic, and Statistical Pattern Recognition, Vol. 8621 of LNCS, Springer Berlin Heidelberg, 2014, pp. 434–443. doi:10.1007/978-3-662-44415-3{_}44.

[67] Z. Wang, A. Bovik, H. Sheikh, E. Simoncelli, Image quality assessment: from error visibility to structural similarity, IEEE Trans. on Image Processing 13 (4) (2004) 600–612. doi:10.1109/TIP.2003.819861.